CLAIMS

What is claimed is:

1. A device to operate as an intermediary node on a network, the device comprising:

a user interface to enable a user to specify a set of forwarding rules for forwarding requests on the network;

a database to store the set of forwarding rules;

a request processing unit to receive a request from a client; and

a rule evaluator to evaluate the set of forwarding rules to identify a rule in the set of forwarding rules which applies to the request, such that the request processing unit attempts to forward the request according to said rule.

2. A device as recited in claim 1, wherein the device operates within a defined forwarding hierarchy, and the user may specify one or more of the forwarding rules to indicate a manner of forwarding the request within the forwarding hierarchy.

3. A device as recited in claim 2, wherein the forwarding hierarchy is a cache hierarchy.

4. A device as recited in claim 3, further comprising a cache to store content requestable by a client on the network, wherein the request processing unit

forwards the request only in the event of a cache miss, wherein the request is a request for content on the network.

5. A device as recited in claim 1, further comprising a rule engine to determine an availability of a host indicated in said rule and to select the host as a forwarding destination for the request if the host is available, the rule engine further to indicate the host to the request processing unit if the host is available to cause the request processing unit to forward the request to the host.

6. A device as recited in claim 1, wherein the rule evaluator identifies the rule which applies to the request by determining that a condition in the rule involving a variable in the request is satisfied.

7. A device as recited in claim 1, wherein, if the request processor is unable to forward the request according to said rule, the rule evaluator resumes evaluating the set of forwarding rules to identify another rule corresponding to the request.

8. A device as recited in claim 1, wherein said rule comprises a plurality of destinations, and wherein the rule engine selects a destination from among the plurality of destinations as a forwarding destination for the request, based on at least one delivery factor included in the rule.

9. A device as recited in claim 8, wherein the at least one delivery factor

comprises a specified distribution method for the request.

10. A device as recited in claim 8, wherein the at least one delivery factor comprises an indication of a current load on at least one of the destinations.

11. A device as recited in claim 8, wherein the at least one delivery factor comprises a weighting of the plurality of destinations indicating a preferred distribution of forwarding requests between the plurality of destinations.

12. A device as recited in claim 1, wherein the user interface is further to enable the user to modify the set of forwarding rules.

13. A device as recited in claim 1, wherein the user interface is further to enable the user to specify a sequence in which the rules of the set of forwarding rules are evaluated in response to the request.

14. A device as recited in claim 1, wherein the user interface is further to enable the user to selectively enable or disable a default forwarding rule.

15. A device to operate as an intermediary node on a network, the device comprising:

a processor;

a network interface to allow the device to communicate on the network;

and

a storage facility to store program code for execution by the processor to cause the device to

provide a user interface to enable a user to specify a set of forwarding rules,

receive a request from a client,

evaluate the set of forwarding rules to identify a rule in the set of forwarding rules which should be applied to the request, and

forward the request on the network according to said rule.

16. A device as recited in claim 15, wherein the device operates within a defined forwarding hierarchy, and the user may specify one or more of the forwarding rules to indicate a manner of forwarding the request within the forwarding hierarchy.

17. A device as recited in claim 16, wherein the forwarding hierarchy is a cache hierarchy.

18. A device as recited in claim 17, further comprising a cache to store content requestable by a client on the network, wherein the device forwards the request only in the event of a cache miss at the device.

19. A device as recited in claim 15, wherein the request is a request for content

on the network.

20. A device as recited in claim 15, wherein the program code is further to cause the device, when executed by the processor, to determine an availability of a host indicated in said rule and to select the host as a forwarding destination for the request if the host is available.

21. A device as recited in claim 15, wherein the program code is further to cause the device, when executed by the processor, to identify the which should be applied in response to the request by identifying a correspondence between a variable in the request and a variable in the rule.

22. A device as recited in claim 15, wherein the user interface is further to enable the user to modify the set of forwarding rules.

23. An intermediary network node comprising:

means for receiving a request for content on a network; and

means for determining a forwarding destination for the request in a defined forwarding hierarchy, by applying a set of user-specified forwarding rules to the request; and

means for forwarding the request according to the determined forwarding destination.

24. An intermediary network node as recited in claim 23, wherein the forwarding hierarchy is a cache hierarchy.

25. An intermediary network node as recited in claim 24, wherein the intermediary node is a network caching device, and wherein said determining and said forwarding are performed only if the request results in a cache miss at the intermediary network node.

26. A caching device to operate within a cache hierarchy on a network, the caching device comprising:

a cache to store content requestable by a client on the network;

a request processing unit to receive a request for content from the client, and to forward the request on the network based on a set of forwarding rules in the event of a cache miss;

a user interface to enable a user to specify the set of forwarding rules, such that the user may specify one or more forwarding rules to indicate a host in the cache hierarchy as a destination for a corresponding request;

a database to store the set of forwarding rules;

a rule evaluator to evaluate the set of forwarding rules in response to the cache miss, to identify a rule in the set of forwarding rules which applies to the request; and

a rule engine to determine an availability of a host indicated in the rule, if any, and to select the host as a forwarding destination for the request if the host is available, the rule engine further to indicate the host to the request processing unit if the host is available to cause the request processing unit to forward the request to the host.

27. An intermediary network node as recited in claim 26, wherein the rule evaluator identifies the rule which applies to the request by determining that a condition in the rule is satisfied.

28. An intermediary network node as recited in claim 27, wherein the condition involves a variable in the request.

29. An intermediary network node as recited in claim 26, wherein said rule comprises a plurality of destinations, and wherein the rule engine selects said destination from among the plurality of destinations, based on the availability of each of the plurality of destinations.

30. An intermediary network node as recited in claim 26, wherein the user interface is further to enable the user to modify the set of forwarding rules.

31. An intermediary network node as recited in claim 26, further comprising a rule encoder to encode into a uniform syntax forwarding rules specified by the

user.

32. A network caching device to operate within a defined cache hierarchy on a network, the caching device comprising:

a cache to store content from an origin server on the network;

an application to receive a request for content from a client via the network, and to forward the request on the network based on a set of forwarding rules in the event of a cache miss;

a user interface to enable a user to specify and modify the set of forwarding rules;

a rule encoder to encode into a uniform syntax forwarding rules specified by the user;

a rules database to store the encoded forwarding rules;

a rule evaluator to evaluate the set of forwarding rules sequentially in response to the cache miss, to identify a rule in the set of forwarding rules which applies to the request, by identifying a correspondence between a variable in the request and a variable in the rule, the rule specifying a host within the cache hierarchy as a forwarding destination for the request; and

a rule engine to determine an availability of the host and to select the host as said forwarding destination for the request if the host is available, the rule engine further to indicate the host to the application layer if the host is available to cause the application layer to attempt to establish a connection with the host,

such that the application layer forwards the request to the host upon successfully establishing the connection.

33. A network caching device as recited in claim 32, wherein said rule comprises a plurality of destinations, and wherein the rule engine selects said destination from among the plurality of destinations, based on the availability of each of the plurality of destinations.

34. A method comprising:

receiving, at an intermediary network node, a request for content on a network;

determining, in the intermediary network node, a forwarding destination in a defined forwarding hierarchy, by applying a set of user-specified forwarding rules to the request; and

forwarding the request according to the determined forwarding destination.

35. A method as recited in claim 34, wherein the forwarding hierarchy is a cache hierarchy.

36. A method as recited in claim 35, wherein the intermediary node is a network caching device, and wherein said determining and said forwarding are performed only if the request results in a cache miss at the intermediary network

node.

37. A method of operating a caching device in a cache hierarchy on a network, the method comprising:

caching content on the network;

providing a user interface to enable a user to specify a set of forwarding rules;

storing the set of forwarding rules;

receiving a request from the client;

evaluating the set of forwarding rules if the request produces a cache miss, to identify a rule in the set of forwarding rules that applies to the request;

determining an availability of a host indicated in the rule;

attempting to establish a connection to the host if the host is available; and

forwarding the request to the host.

38. A method as recited in claim 37, further comprising resuming said evaluating to identify another rule having a correspondence to the request, if said attempting to establish the connection is unsuccessful.

39. A method as recited in claim 37, further comprising, if said attempting to establish the connection is unsuccessful:

determining whether a second available host is indicated in the rule , and

if so,

attempting to establish a connection to the second available host; and

forwarding the request to the second available host.

40. A method as recited in claim 37, wherein said providing the user interface comprises providing the user interface such that the user may specify one or more forwarding rules to indicate a host in the cache hierarchy as a destination for a corresponding request.

41. A method as recited in claim 37, wherein the rule is determined to apply to the request if a condition in the rule is satisfied by the request.

42. A method as recited in claim 41, wherein the condition is a function of a variable in the request.

43. A method as recited in claim 37, wherein said providing the user interface comprises providing the user interface to enable the user to modify the set of forwarding rules.